arXiv:1411.5838v3 [stat.AP] 17 Jun 2015

A Bayesian Particle Filtering Method For Brain Source Localisation

Xi Chen, Simo Särkkä, Simon Godsill

Signal Processing Group, Dept. of Engineering, University of Cambridge
Department of Biomedical Engineering and Computational Science, Aalto University

Abstract

In this paper, we explore the multiple source localisation problem in the cerebral cortex using magnetoencephalography (MEG) data. We model neural currents as point-wise dipolar sources which dynamically evolve over time, then model dipole dynamics using a probabilistic state space model in which dipole locations are strictly constrained to lie within the cortex. Based on the proposed models, we develop a Bayesian particle filtering algorithm for localisation of both known and unknown numbers of dipoles. The algorithm consists of a region of interest (ROI) estimation step for initial dipole number estimation, a Gibbs multiple particle filter (GMPF) step for individual dipole state estimation, and a selection criterion step for selecting the final estimates. The estimated results from the ROI estimation are used to adaptively adjust particle filter's sample size to reduce the overall computational cost. The proposed models and the algorithm are tested in numerical experiments. Results are compared with existing particle filtering methods. The numerical results show that the proposed methods can achieve improved performance metrics in terms of dipole number estimation and dipole localisation.

Keywords: Bayesian, MEG, Multiple source localisation, Particle filter

1. Introduction

In recent years, the development of non-invasive brain signal measuring techniques such as MEG and electroencephalography (EEG) have seen rapid progress. These techniques are helpful in diagnosis of mental diseases such as epilepsy, Alzheimer's and Parkinson's disease [1, 2]. In non-invasive brain

signal processing, we are particularly interested in the signal generated from the cerebral cortex which is the outer layer of the cerebrum [3, 4]. Cortical activity in different cortical regions (such as somatosensory, visual, motor or auditory cortex) can be elicited by suitable stimuli (such as an image or a piece of song). A single active neuron is too weak to be measured directly, so tens of thousands of synchronously active neurons are needed to produce a measurable brain signal. For modelling purposes, many spatially neighbouring active neurons can be summarized and modelled as a dipolar current source, which can be simply named as a "dipole". The electromagnetic field generated by such a dipolar source is measurable using MEG/EEG devices.

Brain source localisation is fundamentally an ill-posed inverse problem [2, 4]. The main barrier is that there may exist many possible solutions for the same set of data, and hence no unique solution can be obtained in the general case. In this paper, we aim to accurately localise the spatio-temporal brain sources using the electromagnetic signals collected outside the surface of the head, employing physiological constraints and soft prior information to regularise the undetermined problem.

1.1. Related work

Brain source localisation is an active research field where a significant amount of work has been done in the past two decades (see, e.g., [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15] and the references therein).

There are two main types of methods: distributed source approaches, and point-wise dipole approaches [3]. Distributed source methods identify the potential active brain sources that are distributed on a dense grid of fixed locations throughout the whole cerebral cortex (or the whole brain volume if under a looser constraint). Since the number of unknown sources is larger than the number of the M/EEG sensors, mathematical assumptions or constraints are required for an unique solution. Some existing methods include the least squares minimum norm estimation (MNE) [3], dynamic statistical parametric mapping (dSPM) [16], standardized low-resolution electromagnetic tomography (sLORETA) [5], and Kalman filter related approaches [8, 6].

On the other hand, point-wise dipole approaches treat the brain currents as point dipole sources, and estimate the states (this may include dipole location, moment, and orientation) of the point source dipoles. In this type of modelling, the state of each dipole source is treated as a random unknown target. A number of works have been published under this type of modelling; these include multiple signal classification (MUSIC) related approaches [17],

Markov chain Monte Carlo related approaches [18, 7], and sequential Monte Carlo (or particle filtering) related approaches [9, 12, 13, 15, 14].

Among the various methods proposed, Bayesian particle filtering seems one of the most promising methods for tackling the source localisation problem. In this paper, we develop a point-wise dipolar source localisation approach using Bayesian particle filtering.

1.2. Bayesian particle filtering methods for dipole localisation problem

Particle filtering methods have been developed for this application over the last decade. Somersalo et al. [9] applied a sequential importance resampling (SIR) particle filter for the dipole localisation problem using artificial planar/3D geometry. Results of a two-dipole localisation example was shown using an ideal spherical head model. Campi et al. [12] proposed a Rao-Blackwellised particle filter (RBPF) for dipole tracking with single dipole and two dipole examples. It was shown in that work that the RBPF provided better localisation results with lower computational cost than those from a standard particle filter. Sorrentino et al. [13] integrated a random finite set scheme into the particle filter. The method was able to track a time-varying number of dipoles with the maximum dipole number specified in advance.

Recently, Sorrentino et. al. [15] suggested to model the problem using a static dipole setup. The work employed a resample-move particle filter to recursively estimate the dipole moment. Chen et al. [19, 20] integrated an MNE step into a multiple particle filter method to localise an unknown number of dipoles. The estimation of the dipole number relied on both the MNE step and the previous localisation history. Miao et al. [14] also adopted a multiple particle filter method to localise multiple dipoles, using a probability hypothesis density (PHD) filter to perform the estimation of the unknown/time-varying dipole number. The algorithm was implemented and assessed in a real-time field-programmable gate array (FPGA) board. However, it modelled the brain under the ideal spherical head model, which cannot provide a realistic description of the true human brain.

1.3. Our work

In this paper, we propose a Gibbs multiple particle filtering (GMPF) algorithm for the multiple dipole source localisation problem. The work is developed based on our previous work [19, 20]. The contribution of this work is described as follows.

Firstly, a continuous head model which forces the state dynamics to strictly remain on the cerebral cortex is developed. To fit with real world applications, we adopt a 1-layer realistic head model, the Nolte model [21]. Although this head model is quite realistic, the off-the-shelf software implementations of it can only be used to evaluate the model at a discrete set of points (the mesh nodes). For distributed source implementations this is all that is needed. However, in our case we need a smooth manifold which defines the cortex surface and hence the discrete set of points is not enough. For this purpose, we adopt a nearest-neighbor (NN) interpolation method to form an approximate continuous cortical manifold. This allows us to formulate the particle filter state directly in terms of the location on the continuous cortex surface.

Secondly, we develop a particle filtering algorithm by integrating a Gibbs sampling iteration step into a multiple particle filtering (MPF) [19] algorithm. Instead of running each component of the MPF only once at each time step, the GMPF iteratively runs the individual components, conditional on the state of the remaining sources, until the state samples converge. This enables the MPF to iterate to obtain a stable state estimate prior to entering the next time step.

Thirdly, we develop a dipole number dynamic model along with the GMPF method [22, 23] for localisation with an unknown dynamic number of dipoles. The model generates three potential dipole number predictions based on the estimate from the previous time step. All three predictions are examined and their corresponding state estimates are calculated. A selection criterion is then applied to obtain the optimal prediction results in each time step. Although approximate in a Bayesian sense, this approach improves the accuracy in estimating the number of dipoles, and thereby improves the overall localisation performance of GMPF.

Finally, we apply a computationally adaptive scheme to adjust the number of particles and the state transition range at each step of the algorithm run. In order to generate candidate numbers of sources at each time step, we integrate a standard noise normalized MNE method [3] and a spatial clustering method [24] to gain some knowledge on the potential dipolar sources. These prior information are used to evaluate the localisation accuracy. We could then adjust the particle size and the state dynamic space in the next algorithm run.

The remainder of the paper is organized as follows. Section 2 introduces the data modelling procedure. A discrete / continuous head model, a dipole

state transition model, and a dipole number dynamic model are described in this section. The localisation algorithm is proposed in Section 3. Both the models and the algorithms are evaluated in Section 4. Section 5 concludes the article.

2. Data model

We consider a clinical application using an MEG system with $M = 204$ magnetometers – the proposed method can be applied to other M/EEG settings with slight modifications. Here we use the 204-sensor MEG application as an example. All the sensors are placed outside the brain surface to obtain non-invasive measurements. We are interested to infer the neural activities within the brain cortical region. The state space is constrained to lie within the cerebral cortex and is denoted as Ω.

For MEG data, a 1-layer realistic head model is introduced to generate the lead-field matrix (the forward matrix), based on a total of G fixed vertices on the cortex. An NN (nearest neighbour) interpolation method is used to interpolate the locations between these vertices.

As described above, the head model comprises G vertices, $\{\mathbf{g}_1 \cdots \mathbf{g}_\nu \cdots \mathbf{g}_G\}$; and F triangular faces on the surface of the cortex, created assuming a 1-shell Nolte model for MEG. The width of the head model is 136 mm, and the distance between two adjacent vertices varies between 2.3 mm to 8.4 mm. The lead-field matrix $\mathbf{L}$ was generated using the statistical parametric mapping (SPM) software [25]. Although $\mathbf{L}$ provides a relatively accurate approximation for the source distribution in the cortical space, it is discretised artificially to a limited number of fixed-location points. We first introduce the traditional discrete real head model using $\mathbf{L}$. The neural current density is, by contrast, in reality a continuous spatial flow. For this reason, we then propose an interpolated realistic head model for continuous point-wise dipole localisation.

Figure 1 shows the triangulation of the cortex. The blue dots are the pre-defined vertices on the cerebral cortex, and the 5 coloured small areas are example sub-planes that represent the individual triangular faces on the cortex. In order to better fit real world applications, we strictly enforce that the trajectory of a point-wise dipolar source lies within the modelled cerebral cortex. Each individual dipolar source may only move within a single triangular cortical region defined by the fixed vertices and the triangular

faces. Thus we model each dipole as semi-static within a small spatial volume for the whole observation interval.

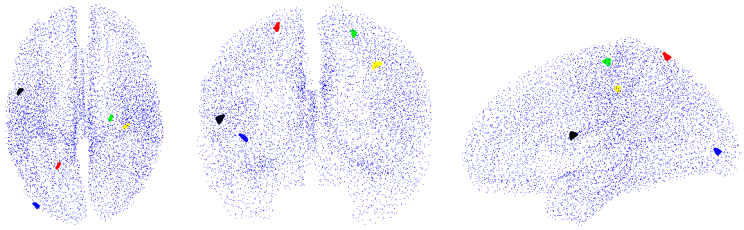

(a) (b) (c)

Figure 1: The vertices (the blue dots) and the individual triangular faces (the five small coloured regions) in the real head model.

For a point-wise dipolar source, we define its state as a vector $\mathbf{X}_k \in \mathbb{R}^{(6N)\times 1}$, where k is the time instant and N_k is the number of dipoles at that time. Since the time resolution of MEG is in the millisecond scale, the unit of the time step in this paper is set as one millisecond. We define the matrix containing the joint state of all dipoles:

$$\mathbf{X}_k = [\mathbf{x}_k^1 \cdots \mathbf{x}_k^n \cdots \mathbf{x}_k^{N_k}]^T. \tag{1}$$

Here, each $\mathbf{x}_k^n$ is a state vector for a single dipole, defined as $\mathbf{x}_k^n = [\mathbf{r}_k^n, \mathbf{q}_k^n]^T$, where $\mathbf{r}_k^n$ is the 3D location and $\mathbf{q}_k^n$ is the dipole moment (the dipole amplitude with orientation). The dipole orientation is set as normal to the cortical subplane surface, so we can usually take $\mathbf{q}_k^n = q_k^n$ to be just the scalar amplitude of the dipole once the cortical geometry is specified by the head model.

A general measurement model that describes the relationship between the MEG measurement and the dipole states is defined as:

$$\mathbf{Y}_k = \mathbf{H}(\mathbf{r}_k)\mathbf{q}_k + \zeta_k, \tag{2}$$

where $\mathbf{Y}_k \in \mathbb{R}^{M\times 1}$ is the measurement vector at time k, $\mathbf{Y}_k = [\mathbf{y}_k^1 \cdots \mathbf{y}_k^m \cdots \mathbf{y}_k^M]^T$. $\mathbf{H}(\cdot)$ represents a linear measurement matrix. $\mathbf{H}(\mathbf{r}_k)\mathbf{q}_k$ is the general measurement model function, $\mathbf{r}_k$ and $\mathbf{q}_k$ denote the vectors of source locations and amplitudes respectively. ζ_k is the measurement noise vector, which is assumed to contain independent Gaussian random variables with zero mean and variance σ_ζ^2.

2.1. Discrete head model

Rather than recomputing, expensively, the general model expression above for each new source location, a discretised head model is generated and interpolated in order to approximate the model on the fine scale of the actual source locations. In order to achieve this, we first compute the lead field matrix corresponding to *every* possible discretised grid location $\{\mathbf{g}_1 \cdots \mathbf{g}_\nu \cdots \mathbf{g}_G\}$. Responses due to arbitrary off-grid source locations are then computed using a special NN interpolation scheme detailed below.

For the discrete head model, the measurement function is simply the discretised lead-field matrix $\mathbf{L}$ [3], an $M \times G$ matrix representing the linear relation between dipole sources at all possible discrete grid locations and the measurements. The model can be rewritten, assuming once again a linear structure, as:

$$\mathbf{Y}_k = \mathbf{L}\mathbf{X}_k^{\mathbf{L}} + \zeta_k^{\mathbf{L}}, \tag{3}$$

where $\mathbf{X}_k^{\mathbf{L}} \in \mathbb{R}^{G\times 1}$ is the vector of amplitudes at the G fixed location vertices in the cortex, $\mathbf{X}_k^{\mathbf{L}} = [q_{k,\mathbf{g}_1} \cdots q_{k,\mathbf{g}_\nu} \cdots q_{k,\mathbf{g}_G}]^T$. Here, for the discrete model, the set of dipole locations $\mathbf{r}_k^n$ are pre-specified from the fixed spatial grid $[\mathbf{g}_1 \cdots \mathbf{g}_\nu \cdots \mathbf{g}_G]^T$.

2.2. Continuous head model

We apply a simple NN interpolation method [26] between the G fixed location anchor vertices in order to obtain the approximated continuous head model. As shown in Figure 2, for the three closest neighbouring anchor vertices to source location $\mathbf{r}_k$ (denoted as $\mathbf{g}_1$, $\mathbf{g}_2$ and $\mathbf{g}_3$ in the example), we obtain $\mathbf{Y}^{\mathbf{g}_1}$, $\mathbf{Y}^{\mathbf{g}_2}$ and $\mathbf{Y}^{\mathbf{g}_3}$. $\mathbf{Y}^{\mathbf{g}_\nu}$ represents the unit noiseless response measured by MEG when we place a unit dipole at the νth anchor vertex $\mathbf{g}_\nu$, computed using [25]. The area of the triangular region varies depending on the distance between the three anchor vertices, with the average area at around 7 mm^2. The triangular cortical region is sufficiently small that it is reasonable to treat it as a flat sub-plane in which all the interpolated points have the same orientation. In Figure 2, the orientation that is normal to the sub-plane is denoted as e_p. e_1, e_2 and e_3 are the orientations of the three anchor vertices respectively when there is a unit dipole placed at each vertex. We define θ as the angle between e_p and the orientation of the anchor vertex. We have then θ_1, θ_2 and θ_3 for vertices $\mathbf{g}_1$, $\mathbf{g}_2$ and $\mathbf{g}_3$, respectively.

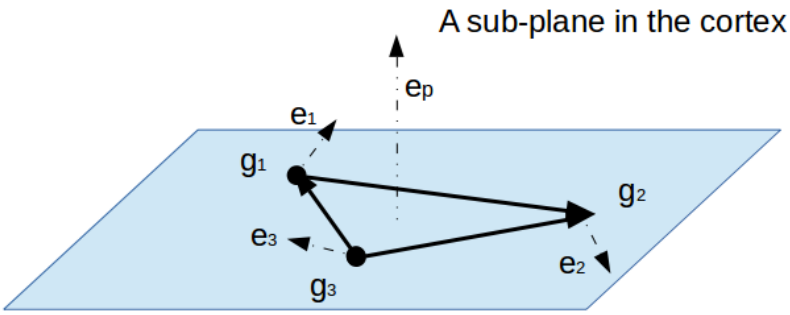

Figure 2: The nearest-neighbour interpolation for a three neighbouring vertices example.

For an nth dipole with location $\mathbf{r}_k^n$ at k, we compute its unit response as:

$$\mathbf{Y}^{\mathbf{r}_k^n} = (1-\phi-\varphi)\widetilde{\mathbf{Y}}^{\mathbf{g}_3} + \phi\widetilde{\mathbf{Y}}^{\mathbf{g}_2} + \varphi\widetilde{\mathbf{Y}}^{\mathbf{g}_1}, \tag{4}$$

where $\widetilde{\mathbf{Y}}^{\mathbf{g}_\nu} = \cos\theta_\nu \mathbf{Y}^{\mathbf{g}_\nu}$ is the unit response after orientation mapping, ϕ and φ are interpolation coefficients describing the relation between the location of a dipole and the three anchor vertices, computed as follows.

To obtain the relations between ϕ, φ and the 3D locations $\mathbf{r}_k^n$, $\mathbf{g}_1$, $\mathbf{g}_2$, and $\mathbf{g}_3$, we define $\alpha = \mathbf{g}_1 - \mathbf{g}_3$, $\beta = \mathbf{g}_2 - \mathbf{g}_3$ and $\gamma = \mathbf{r}_k^n - \mathbf{g}_3$. We have [26]:

$$\phi = \frac{(\alpha\cdot\beta)(\gamma\cdot\beta) - (\beta\cdot\beta)(\gamma\cdot\alpha)}{(\alpha\cdot\beta)^2 - (\alpha\cdot\alpha)(\beta\cdot\beta)}, \tag{5}$$

$$\varphi = \frac{(\alpha\cdot\beta)(\gamma\cdot\alpha) - (\alpha\cdot\alpha)(\gamma\cdot\beta)}{(\alpha\cdot\beta)^2 - (\alpha\cdot\alpha)(\beta\cdot\beta)}. \tag{6}$$

We then obtain the final measurement by summing up all of the scaled unit responses multiplied by its amplitude, from different individual dipoles: $\mathbf{Y}_k = \sum_{n=1}^{N} \mathbf{Y}^{\mathbf{r}_k^n} q_k^n \simeq \mathbf{H}(\mathbf{r}_k)\mathbf{q}_k$, as in Equation 2, where the approximation arises as a result of the NN interpolation procedure. The interpolation is executed in the following steps:

- Find the triangular sub-planes where the dipoles $\mathbf{X}_k$ are located in.
- For each individual dipole $\mathbf{x}_k^n$, identify the anchor vertices $\mathbf{g}_\nu$ and compute their corresponding orientations.

- Calculate the angles θ_ν to obtain $\widetilde{\mathbf{Y}}_k^{\mathbf{g}_\nu}$.
- Calculate ϕ and φ using Equation (5) and (6), then compute $\mathbf{Y}^{\mathbf{r}_k^n}$.
- Sum up all $\mathbf{Y}_k^{\mathbf{r}_k^n}$ to obtain the predicted (noiseless) measurement $\mathbf{Y}_k$.

2.3. Individual dipole dynamic model

For an individual dipole which exists from time instant $k-1$ to k, we define the following individual dipole transition model:

$$\mathbf{x}_k^n = f(\mathbf{x}_{k-1}^n, S_\kappa^n, \phi^n, \varphi^n), \tag{7}$$

where $f(\cdot)$ can be a linear or nonlinear function. S_κ is the sub-plane, κ is the face index and n is the dipole index. Each of the F sub-planes in the 3D cortical space can be treated as a two dimensional plane. Thus we adopt a two dimensional random walk model as our transition function $f(\cdot)$. As we have stated above, the location of a dipole is modelled as semi-static on the cerebral cortex. The state space is constrained accurately on the brain cortical surface Ω, thereby we divide the dipole dynamic into two phases: the transition between different triangular faces and the transition within an individual face.

For the transition between different faces, we define $\vartheta(\cdot)$ as the neighbouring sub-plane set, which stores all of the neighbouring triangular faces to the face where the dipole $\mathbf{x}_{k-1}^n$ is located at the previous time step. In practice, this is computed and stored in a lookup table using the grid information provided by the lead-field matrix.

For the transition within a face, we draw values for the coefficients ϕ^n and φ^n, and randomly select a position within the triangular sub-plane, $\phi^n \sim \mathbb{U}[0,1]$ and $\varphi^n \sim \mathbb{U}[0,1]$ with constraint described in Equation (4).

The procedure to perform the dipole transition is as follows:

- For a dipole with state $\mathbf{x}_{k-1}^n$, find the neighbour sub-plane set $\vartheta(\mathbf{x}_{k-1}^n)$.
- We have $F_\kappa^n + 1$ sub-planes including F_κ^n neighbouring sub-planes $S_\kappa^n \in \vartheta(\mathbf{x}_{k-1}^n)$ and the original sub-plane where $\mathbf{x}_{k-1}^n$ located. We randomly select one sub-plane with equal probability $p_\kappa^n = \frac{1}{F_\kappa^n + 1}$.
- Randomly choose a position using ϕ^n and φ^n in the selected sub-plane.
- Calculate $\mathbf{x}_k^n = f(\mathbf{x}_{k-1}^n, S_\kappa^n, \phi^n, \varphi^n)$ to obtain the new dipole state prediction.

F_κ^n is the total number of the identified neighbour faces. In the initialisation step, dipole states are randomly selected from the G fixed location grid points. Since a grid point connects several sub-planes, it is difficult to define which sub-plane it belongs to. In practice, we compute the distance between the grid point and all its neighbour sub-planes. We select the one with the shortest distance to the grid point as its sub-plane.

2.4. Dipole number dynamic model

In our model, we assume that the number of the active dipolar sources does not change dramatically between adjacent time steps. For scenarios with more than one dipole currently active, we can model the dipole number transition process by allowing an individual dipole to appear or disappear at a single time instant.

Given the dipole number N_{k-1} from the previous time instant, the current dipole number N_k can be obtained from the following dipole number dynamic model:

$$p(N_k \mid N_{k-1}) = \begin{cases} p_{k,(+)} & \text{for } N_k = N_{k-1} + 1 \\ p_{k,(0)} & \text{for } N_k = N_{k-1} \\ p_{k,(-)} & \text{for } N_k = N_{k-1} - 1 \end{cases} \tag{8}$$

with $\sum_j p_{k,j} = 1$, where $j = \{(-), (0), (+)\}$. The dynamic probability $p_{k,j}$ is predefined with a consideration of dipole birth-death movement [27]. In general, we set $p_{k,(+)} = p_b$, $p_{k,(-)} = \frac{p_d}{N_{k-1}}$, $p_{k,(0)} = 1 - \frac{p_b - p_d}{N_{k-1}}$. p_b and p_d are the birth-death probability discussed as follows.

2.4.1. Dipole birth-death move

We adopt a simple birth-death move for cases when a new dipole appears or an existing dipole disappears from time $k-1$ to k. We set $p_b = p_d = \frac{p_i}{2}$, where $p_i \sim \mathbb{U}(0, 1)$. For initialisation step or special cases when $N_{k-1} = 0$, we set $p_d = 0$ and $p_b = p_i$.

For the birth process, a randomly selected initial state within the cortical surface Ω is assigned for the new birth dipole. The initialisation is assisted by the results from the probabilistic ROI estimation (detailed in Section 3). The probabilistic ROI estimation provides us knowledge about the signal strength of the G grid points in L. A probabilistic sampling method is used for sample selection regarding the signal strength (explained in Section 3.2). The selected sample point set is defined as Ψ_k. We randomly select one of the grid points in Ψ_k and assign its location to the new birth dipole. Ψ_k is

also used to generate the ROIs by applying a spatial clustering algorithm [24] to the selected point sets.

For the dipole death process, the algorithm randomly selects and deletes one of the existing dipoles. The random selection does not introduce extra bias in the death process due to the large sample number and resampling step in the particle filter. Instead of computing a single result in each time step, the particle filter will draw a number of sample points and compute their weights; samples with lower weights are filtered out to improve the estimation accuracy. Therefore the deleted existing dipoles vary from sample to sample, only samples with a higher posterior density will be used for the final estimates. The particle filtering algorithm will be detailed in Section 3.

3. Bayesian sequential Monte Carlo algorithm

In this section, we develop a Bayesian particle filter algorithm to deal with the localisation with both known and unknown dipole number. The dipole localisation problem is cast as a general multi-target tracking problem in our Bayesian framework.

3.1. Algorithm execution

We first describe the algorithm structure, followed by a detailed description of each component in the proposed algorithm. The algorithm contains three different interactive components: the probabilistic ROI estimation step estimates the number of ROIs $\overline{N}_k$ at time k; $\overline{N}_k$ provides an initial guess of the dipole number for the main algorithm; the GMPF step performs the particle filtering; the selection criterion step selects the optimal number $\hat{N}_k$ and its corresponding state estimates $\hat{\mathbf{X}}_k$ at time k. Figure 3 illustrates the relationship between the components and their corresponding variables. The target number propagation model and the selection criterion scheme are relatively naive and simple when compared to some other existing estimation methods [27]. In this paper we focus on the tracking algorithm performance.

The algorithm executes in the following four steps:

1. A probabilistic ROI estimation step uses the noise normalised MNE method and a probabilistic sampling method to obtain the point set Ψ_k. Therefore we can obtain an estimate for the number of the ROIs $\overline{N}_k$. $\overline{N}_k$ is used for the dipole number initialisation at N_0 at $k = 0$.

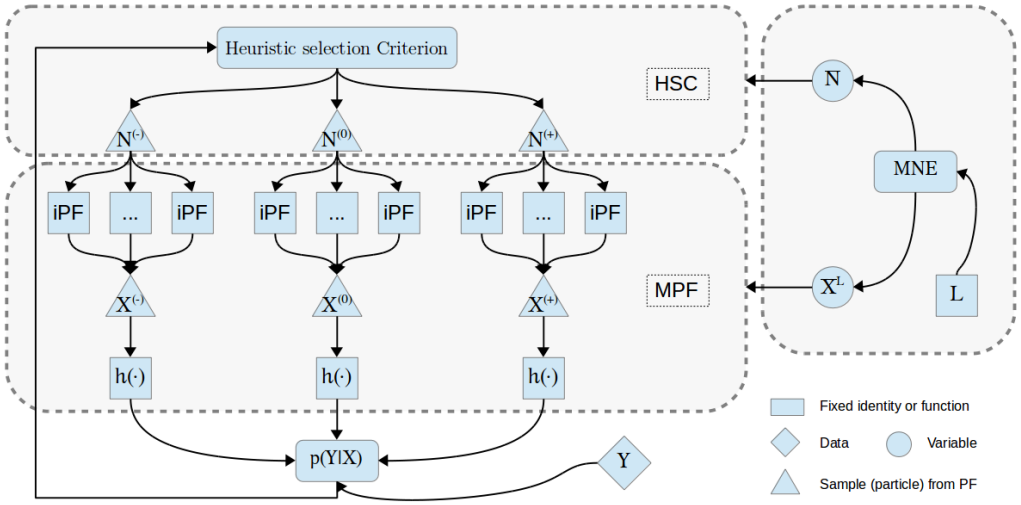

Figure 3: Algorithm execution illustration

2. Given the estimated dipole number $\hat{N}_{k-1}$ from $k-1$, three potential dipole numbers N_k^j (with their corresponding dynamic probability) can be obtained from Equation (8).
3. For each N_k^j, the GMPF algorithm is applied to estimate the dipole state. The algorithm assigns an individual particle filter (iPF) for each of the identified targets. In a single GMPF run, each iPF is executed to generate the estimate $\mathbf{x}_k^n$ for its corresponding dipole. $\mathbf{x}_k^n$ is immediately updated in the dipole state $\mathbf{X}_k^j$ to assist other iPFs. This updating procedure is performed in a Gibbs sampling manner (see Section 3.3.1 for details). We finally obtain the estimated dipole state $\widetilde{\mathbf{X}}_k^j$ for each of the three potential j cases.
4. The selection criterion scheme is then used to compute the posterior probability for each of the three cases. One of the three cases is selected as the final estimate at k. An estimation for both the dipole number $\hat{N}_k$ and the dipole state $\hat{\mathbf{X}}_k$ can then be obtained. We set $N_k = \hat{N}_k$ and $\mathbf{X}_k = \hat{\mathbf{X}}_k$.

In addition, the discrete point set Ψ_k from the MNE is used to adaptively control the number of particles and the state transition range in the particle filtering. In the remainder of this section, we give a more detailed description for each parts.

3.2. Probabilistic ROI estimation

A noise-normalised MNE approach is employed to obtain the signal strength of the fixed grid points. The discrete head model from Equation (3) is $\mathbf{Y}_k = \mathbf{L}\mathbf{X}_k^{\mathbf{L}} + \zeta_k^{\mathbf{L}}$. A typical noise-normalised MNE solution [3] can be obtained as follows:

$$\overline{\mathbf{X}}_k^{\mathbf{L}} = \mathbf{L}^T(\mathbf{L}\mathbf{L}^T + \lambda I)^{-1}\mathbf{Y}_k, \tag{9}$$

where $\overline{\mathbf{X}}_k^{\mathbf{L}}$ is an amplitude estimation vector for the G grid points, $\overline{\mathbf{X}}_k^{\mathbf{L}} = [\overline{q}_{k,\mathbf{g}_1} \cdots \overline{q}_{k,\mathbf{g}_\nu} \cdots \overline{q}_{k,\mathbf{g}_G}]^T$. λ is a noise normalised regularisation parameter.

We do a simple probabilistic sampling with respect to the estimated amplitude of the G points, as follows:

$$\mathbf{g}_{k,\nu} \sim \frac{\overline{q}_{k,\nu}}{\sum_{\nu=1}^{G} \overline{q}_{k,\nu}}. \tag{10}$$

We now have the point set $\Psi_k = \{\mathbf{g}_{k,1}, \cdots, \mathbf{g}_{k,\nu}, \cdots, \mathbf{g}_{k,G}\}$. We employ a hierarchical spatial clustering method [24] to cluster the selected points in Ψ_k with respect to their corresponding geographical positions. We define these clusters as the ROIs. We can obtain from this the number and extent of the ROIs $\Psi_k(n)$, $n = 1, 2, \cdots, \overline{N}_k$.

$\Psi_k(n) = \{\mathbf{g}_{k,1} \cdots \mathbf{g}_{k,\nu^n} \cdots \mathbf{g}_{k,G^n}\}$ is the subset of the grid points for active region n, where G^n is the total number of points in that subset and ν^n is the grid point index. We compute the location of the n_kth ROI by taking the mean of the points in $\Psi_k(n)$ so that the centre of the ROI is defined as:

$$\mathbf{c}_k^n = \frac{1}{G^n} \sum_{\nu^n=1}^{G^n} \mathbf{g}_{k,\nu^n}. \tag{11}$$

The centre location of each ROI will be used in Section 3.5.

3.3. Multiple particle filtering

For each value of N_k^j, the problem reduces to track a fixed number of dipoles. We target the joint filtering distribution for all sources, $p(\mathbf{X}_k^j \mid \mathbf{Y}_{1:k}, N_{1:k-1}, N_k^j)$, where $N_{1:k-1}$ is the previous estimated dipole number up to time $k-1$, and N_k^j is the current value of dipole number for the jth case. The state vector of each GMPF is denoted by $\mathbf{X}_k^j$. We assign a GMPF for each N_k^j; the three GMPFs operate in parallel and are independent of each other.

For each j, there are N_k^j dipoles and therefore N_k^j iPFs are assigned, one for each active source, as shown in Figure 3. The state vector can be expanded as $\mathbf{X}_k^j = \{\mathbf{x}_k^1, \mathbf{x}_k^2, \cdots \mathbf{x}_k^{n_k^j}, \cdots \mathbf{x}_k^{N_k^j}\}$, where $\mathbf{x}_k^{n_k^j}$ denotes the state of the n_k^jth individual dipole out of N_k^j. To simplify the notation, rewrite the state $\mathbf{x}_k^{n_k^j}$ as $\mathbf{x}_{n,k}^j$.

Now, defining $\psi_k^j = \{N_{1:k-1}, N_k^j\}$ as the trajectory of the previous and the current active source numbers, the joint target distribution $p(\mathbf{X}_k^j \mid \mathbf{Y}_{1:k}, \psi_k^j)$ can be expressed as:

$$p(\mathbf{X}_k^j \mid \mathbf{Y}_{1:k}, \psi_k^j) = \frac{p(\mathbf{Y}_{1:k} \mid \mathbf{X}_k^j, \psi_k^j) p(\mathbf{X}_k^j \mid \psi_k^j)}{p(\mathbf{Y}_{1:k} \mid \psi_k^j)}. \tag{12}$$

As ψ_k^j is known in the current time step and the dipoles are assumed to be independent of each other *a priori*, we have:

$$p(\mathbf{X}_k^j \mid \mathbf{Y}_{1:k}, \psi_k^j) \propto p(\mathbf{Y}_{1:k} \mid \mathbf{x}_{1,k}^j, \mathbf{x}_{2,k}^j, \cdots, \mathbf{x}_{N,k}^j, \psi_k^j) \prod_{n=1}^{N} p(\mathbf{x}_{n,k}^j \mid \psi_k^j), \tag{13}$$

where $p(\mathbf{x}_{n,k}^j \mid \psi_k^j)$ is the predictive distribution for each independent source; the states are uniformly drawn from the whole state space Ω.

In order to draw samples from the joint target posterior distribution, we consider a Gibbs sampling structure [28] to iteratively update the desired dipole states at each time step, in which a conditional particle filter for each active source approximates the required full conditional draw required for a Gibbs sampler. The basic idea is very simple: at time $k-1$ we assume that particles are available from from the joint target $p(\mathbf{X}_{k-1} \mid \mathbf{Y}_{1:k-1}, \psi_{k-1})$. In order to run this forward one time step, we initialise the state $\mathbf{X}_k$ to some arbitrary initial values. We then run conditional particle filters for each target in turn. Each particle filter is desiged to target the conditional filtering distribution $p(\mathbf{x}_{n,k} \mid \mathbf{x}_{-n,k}, \mathbf{Y}_{1:k}, \psi_k)$, where $\mathbf{x}_{-n,k}$ denotes the state $\mathbf{X}_k$ with source number n removed. A single sample is randomly selected from the particle approximation to $p(\mathbf{x}_{n,k} \mid \mathbf{x}_{-n,k}, \mathbf{Y}_{1:k}, \psi_k)$ and this is substituted into the state vector $\mathbf{X}_k$ at the nth source position. One iteration of the Gibbs sampler comprises a complete sweep through all of the particle conditional distributions, and convergence will occur after some large number of iterations (in practice though we implemented just a few iterations). In this way we aim to overcome the approximation induced by the standard MPF, which

accounts only approximately for the statistical coupling between sources, at the expense of an iterative procedure at each time step. An interesting modification of our approach would be to incorporate pseudo-marginal sampling ideas into the Gibbs sampler, which would in addition correct the approximation of the particle filter to the conditional filtering distribution. This addition is left as a future topic for exploration.

We now specify in detail the steps required for the GMPF.

3.3.1. Gibbs sampling updating step for conditional posterior distribution

In each time step, the iPFs are executed from the first iPF to the N_k^jth iPF sequentially. The updating procedure acts in a similar way to that in a standard Gibbs sampler: once a new state $\mathbf{x}_{n,k}^j$ is generated, it is immediately used to assist other iPFs by updating the corresponding conditional posterior distribution. In order to obtain a good set of samples from the joint filtering distribution this Gibbs sampler should run for a number of iterations at each time step. We denote the iteration number by the iteration indices $l' = \{1, 2 \cdots l\}$.

The conditional posterior distribution of the individual dipole state which is approximated by each iPF can be written as $p(\mathbf{x}_{n,k}^{j(l')} \mid \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})$. The iPF state $\mathbf{x}_{n,k}^{j(1)}$ is initialised to $\mathbf{x}_{n,k-1}^{(l)}$ from the previous time step after the selection criterion step. $\mathbf{x}_{-n,k}^{j(l')}$ describes a state vector excluding the state $\mathbf{x}_{n,k}^j$ at the l'th (when $l' \geq 2$) Gibbs iteration:

$$\mathbf{x}_{-n,k}^{j(l')} = \{\mathbf{x}_{1,k}^{j(l')}, \cdots \mathbf{x}_{n-1,k}^{j(l')}, \mathbf{x}_{n+1,k}^{j(l'-1)} \cdots \mathbf{x}_{N,k}^{j(l'-1)}\}. \tag{14}$$

$\mathbf{x}_{-n,k}^{j(l')}$ is used in the estimation of the nth iPF, its first $n-1$ elements are the estimates updated in the current Gibbs iteration while the other elements are the estimates from the previous Gibbs iteration. Once a new state $\mathbf{x}_{n,k}^{j(l')}$ is generated, we update the corresponding $\mathbf{x}_{-n,k}^{j(l')}$ and use it in the $(n+1)$th iPF. The basic scheme of the Gibbs iteration is described as follows:

- Generate $\mathbf{x}_{1,k}^{j(l')}$ from $p(\mathbf{x}_{1,k}^j \mid \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-1,k}^{j(l')})$.
- Generate $\mathbf{x}_{2,k}^{j(l')}$ from $p(\mathbf{x}_{2,k}^j \mid \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-2,k}^{j(l')})$.

  $\vdots$
- Generate $\mathbf{x}_{N,k}^{j(l')}$ from $p(\mathbf{x}_{N,k}^j \mid \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-N,k}^{j(l')})$.

Therefore, the final iteration gives $\mathbf{X}_k^{j(l)} = \{\mathbf{x}_{1,k}^{j(l)}, \mathbf{x}_{2,k}^{j(l)}, \cdots \mathbf{x}_{n,k}^{j(l)}, \cdots \mathbf{x}_{N,k}^{j(l)}\}$, where $\mathbf{x}_{n,k}^{j(l)}$ denotes the sampled value of a dipole state at time step k for the jth case, obtained after l iterations of Gibbs update. After the selection criterion step for each algorithm run, one of the three j cases is selected, we will have $\hat{\mathbf{X}}_k$ where the notation j is eliminated. This is important because we use some variables at $k-1$ in our Bayesian inference. For example, the term $\psi_{k-1} = \{N_{1:k-2}, N_{k-1}\}$ and $\mathbf{x}_{-n,k-1} = \{\mathbf{x}_{1,k-1}, \mathbf{x}_{2,k-1}, \mathbf{x}_{N,k-1}\}$.

The Gibbs sampling iteration enables the algorithm to get a more accurate estimate in each particle filtering step, particularly when the sources are spatially close and hence not independent in their joint posterior. This raises several issues such as the Gibbs iteration number, convergence analysis, and the computational load considerations. In theory, many iterations would be required to guarantee convergence, but here we only operate a few iterations because of the high computational burden. The intuition here is that usually the dipoles are well separated and quite independent; hence the Gibbs sampling should be approximately converged within a few iterations.

3.3.2. Individual particle filtering

Since ψ_k^j and $\mathbf{x}_{-n,k}^{j(l')}$ are available terms to each iPF, we only need to sample the unknown state $\mathbf{x}_{n,k}^{j(l')}$ in each iPF. The conditional posterior distribution can be rewritten as $p(\mathbf{x}_{n,k}^j \mid \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})$. We define I_k as the number of particles for each iPF at time k. A particular sample in an iPF is denoted as $\mathbf{x}_{n,k}^{ij}$, which is updated in each of the l'th Gibbs iteration.

The conditional posterior distribution required for Gibbs sampling can be expanded in two steps: an updating step and a prediction step. The (conditional) updating step can be expressed as:

$$p(\mathbf{x}_{n,k}^j \mid \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')}) = \frac{p(\mathbf{Y}_k \mid \mathbf{x}_{n,k}^j, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')}) p(\mathbf{x}_{n,k}^j \mid \mathbf{Y}_{1:k-1}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})}{p(\mathbf{Y}_k | \mathbf{Y}_{1:k-1}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})}, \tag{15}$$

where

$$\begin{aligned} &p(\mathbf{x}_{n,k}^j \mid \mathbf{Y}_{1:k-1}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')}) \\ &= \int p(\mathbf{x}_{n,k}^j \mid \mathbf{x}_{n,k-1}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')}) p(\mathbf{x}_{n,k-1} \mid \mathbf{Y}_{1:k-1}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')}) d\mathbf{x}_{n,k-1} \end{aligned}$$

is the corresponding prediction step. The conditional posterior for $\mathbf{x}_{n,k-1}$ at $k-1$ is $p(\mathbf{x}_{n,k-1} \mid \mathbf{Y}_{1:k-1}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})$, containing parameters ψ_k^j and $\mathbf{x}_{-n,k}^{j(l')}$ which are known prior to the iPF run. The conditional posterior distribution $p(\mathbf{x}_{n,k}^j \mid \mathbf{Y}_{1:k}, \psi_{k+1}^j, \mathbf{x}_{-n,k+1}^{j(l')}) \approx p(\mathbf{x}_{n,k}^j \mid \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})$, which forms a complete recursive form in the Bayesian model. This is because the dipole sources are a priori independent of each other, conditioning on ψ_{k+1}^j and $\mathbf{x}_{-n,k+1}^{j(l')}$ does not affect the proposed conditional posterior distribution. Here we simply take the weighted samples from the nth iPF, building in the independence approximation

$$p(\mathbf{X}_k^j \mid \mathbf{Y}_{1:k}, \psi_{k+1}^j) \approx \prod_{n=1}^{N_k^j} p(\mathbf{x}_{n,k}^j \mid \mathbf{Y}_{1:k}, \psi_{k+1}^j, \mathbf{x}_{-n,k+1}^{j(l')}). \tag{16}$$

In these equations we have also built in the assumptions that the sources are a priori independent at each time point, and that the observations at time k are conditionally independent of the states prior to time k. We choose an importance density $q(\mathbf{x}_{n,k}^j \mid \mathbf{x}_{n,k-1}, \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})$ and then obtain the weight as

$$w_{n,k}^{ij} \propto w_{n,k-1}^i \frac{p(\mathbf{Y}_k \mid \mathbf{x}_{n,k}^{ij}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')}) p(\mathbf{x}_{n,k}^{ij} \mid \mathbf{Y}_{1:k-1}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})}{q(\mathbf{x}_{n,k}^{ij} \mid \mathbf{x}_{n,k-1}^i, \mathbf{Y}_{1:k}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})}, \tag{17}$$

where $\mathbf{x}_{n,k-1}^i$ denotes the ith particle sample after the selection criterion step at $k-1$. We choose the prior as the importance distribution, so that we have the following simplification, the standard bootstrap filter [29]:

$$w_{n,k}^{ij} \propto w_{n,k-1}^i p(\mathbf{Y}_k \mid \mathbf{x}_{n,k}^{ij}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')}). \tag{18}$$

We normalise $w_{n,k}^{ij}$ to obtain $\widetilde{w}_{n,k}^{ij}$. We adopt a residual resampling step [30] to avoid the so-called degeneracy problem [31]. See Algorithm 1 for the detailed description of this iPF.

For the particle weight of the iPF in the birth move (a new target appears), the samples are drawn uniformly from the whole state space Ω, and assigned equal weight $\frac{1}{I_k}$. For the death move (an existing target disappears), the corresponding particles of the selected target are deleted.

As stated in Section 3.3.1, $\mathbf{x}_{-n,k}^{j(l')}$ is then updated using the result from the current iPF run. Since it is impossible to obtain the ground-truth state

Algorithm 1: Individual particle filter

`// At time` k `for the` n_k^j`th dipole of a total` N_k^j `dipoles, at the` l'`th Gibbs sampling iteration with` I_k `particles`

for $i = 1, \ldots, I_k$ **do**

`// Prediction`

- Draw samples $\mathbf{x}_{n,k}^{ij} \sim p(\mathbf{x}_{n,k}^{j} | \mathbf{x}_{n,k-1}^{i}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})$.
- Compute weights: $w_{n,k}^{ij} \propto w_{n,k-1}^{i} p(\mathbf{Y}_k | \mathbf{x}_{n,k}^{ij}, \psi_k^j, \mathbf{x}_{-n,k}^{j(l')})$.

end

- Normalise weights $\tilde{w}_{n,k}^{ij} = \frac{w_{n,k}^{ij}}{\sum_i w_{n,k}^{ij}}$.

`// Resample`

$\{\mathbf{x}_{n,k}^{ij}, \tilde{w}_{n,k}^{ij}\}_{i=1}^{I_k}$ to $\{\mathbf{x}_{n,k}^{(i'j)}, \frac{1}{I_k}, i'\}_{i'=1}^{I_k}$.

`// Gibbs iteration choice`

- Update the n_k^jth state estimate $\mathbf{x}_{n,k}^{j}$ by random selection from $\mathbf{x}_{n,k}^{(i'j)}$.
- Assign $\mathbf{x}_{n,k}^{j(l')} = \mathbf{x}_{n,k}^{j}$ in the Gibbs iteration.

$\mathbf{x}_{n,k}^{j}$ for each n_k^jth dipole source directly, we randomly pick up the state estimate from the resampled state $\mathbf{x}_{n,k}^{ij}$. When the particle weights at k are not available, we randomly select a sample from $\mathbf{x}_{n,k}^{i}$ according to the weight $\tilde{w}_{n,k-1}^{ij}$ from the previous time step. Once the current particle weights are obtained, we can obtain the estimate $\mathbf{x}_{n,k}^{j}$ by randomly selecting a sample from the resampled particles $\{\mathbf{x}_{n,k}^{(i'j)}\}$ with an equal probability $p^{(i)} = \frac{1}{I_k}$, where i' represents the particle index after resampling.

The main body of the proposed algorithm is described in Algorithm 2.

To sum up, we modify the original MPF algorithm in the following three ways:

(1) We integrate an iterated Gibbs sampling procedure to generate the individual state estimate in each iPF. This enables us to obtain more reliable estimations in each $\mathbf{x}_{-n,k}^{j}$ assisted iPF run. The number of iterations is controlled by the parameter l.

(2) Instead of dividing the state space into several subspaces, samples of each iPF are drawn from the same state space. In the dipole initialisation step, the samples are drawn from the G vertices; they are propagated using the individual dipole dynamic model in Ω.

(3) Rather than using a weighted mean method, $\mathbf{x}_{n,k}^{j}$ in the nth iPF is selected randomly from all the samples $\mathbf{x}_{n,k}^{ij}$, where i is the sample index. In practice, we randomly pick up one of the particle filter samples $\mathbf{x}_{n,k}^{ij}$ with the equal probability $p^{(i)} = \frac{1}{I_k}$.

3.4. Selection criterion scheme

We then obtain three N_k^j candidates and the corresponding states $\widetilde{\mathbf{X}}_k^j$ from the GMPF. We apply the selection criterion scheme to find the optimal pair amongst all the available estimates. We can obtain the posterior probability as

$$p(N_k^j \mid \mathbf{Y}_{1:k}, \hat{N}_{1:k-1}) \propto p(\mathbf{Y}_k \mid \mathbf{Y}_{1:k-1}, \psi_k^j) p(N_k^j \mid \hat{N}_{k-1}). \tag{19}$$

The estimate is:

$$\hat{N}_k = \arg\max p(\mathbf{Y}_k \mid \mathbf{Y}_{1:k-1}, \psi_k^j) p(N_k^j \mid \hat{N}_{k-1}), \tag{20}$$

where $p(\mathbf{Y}_k \mid \mathbf{Y}_{1:k-1}, \psi_k^j) \approx \frac{1}{N_k^j} \sum_{i=1}^{I_k} \prod_{n=1}^{N_k^j} w_{n,k}^{ij}$. According to [31], and assuming that the source posterior factorizes over n (i.e., the sources are independent), we can then obtain $\hat{N}_k$ by selecting the N_k^j with the highest probability. Then obtaining $\hat{\mathbf{X}}_k$ from the corresponding posterior mean $\widetilde{\mathbf{X}}_k^j$.

3.5. Adaptive filtering

As shown in Figure 3, the discrete amplitude matrix $\overline{\mathbf{X}}_k^{\mathbf{L}}$ from MNE is used in GMPF to assist the sampling procedure in every iPF run. The identified ROIs and their point sub-set $\Psi_k(n)$ are used to control the particle number and particle transition range. In practice, we calculate the localisation root mean squared error (RMSE) between the centre of each ROI and the dipole state estimation $\mathbf{X}_k$ at each time step k. The localisation RMSE e_k can be obtained from:

$$e_k = D(\Psi_k(n), \mathbf{X}_k), \tag{21}$$

where function $D(\cdot)$ computes the localisation RMSE between all elements in $\Psi_k(n)$ and $\mathbf{X}_k$. The centre points of the ROIs $\mathbf{c}_k^n$ are compared with the localisation estimation$\mathbf{X}_k$. We compute the spatial distance $||\mathbf{X}_k - \mathbf{c}_k^n||$ between each of the the individual dipole state and the ROIs. We then obtain an $N_k \times \mathbf{g}_k^n$ matrix C_k that contains all pairs of the localisation RMSE. We

Algorithm 2: Bayesian multiple dipole localisation algorithm

Initialisation at $k = 0$: compute $\overline{\mathbf{X}}_0^{\mathbf{L}}$, assign $N_0 = \overline{N}_0$.
Randomly draw $\mathbf{X}_0$ from Ω for each of the three cases.
for $k = 1, \ldots, K$ **do**
  `// Probabilistic ROI estimation`
  Compute $\overline{\mathbf{X}}_k^{\mathbf{L}}$ to obtain $\overline{N}_k$ and Ψ_k (see Section 3.2).
  for $j = (-), (0), (+)$ **do**
    See Equation (8).
    Set $\mathbf{X}_k^j = \hat{\mathbf{X}}_{k-1}$.
    if $j = (+)$ **then**
      `// Birth move`
      Uniformly draw I_k particles from Ω for the new dipole.
      Calculate its state and append it to $\mathbf{X}_k^j$.
    else
      `//` $j = (-)$ `Death move`
      Randomly select a dipole estimate from $\mathbf{X}_k^j$. Delete the selected state and its corresponding I_k particles.
    end
    Initialise Gibbs step $\mathbf{X}_k^{j(1)} = \mathbf{X}_k^j$.
    for $l' = 1, 2 \ldots l$ **do**
      `// Gibbs iteration`
      for $n = 1, \ldots, N_k^j$ **do**
        `// iPF`
        Follow Algorithm 1 to obtain $\mathbf{x}_{n,k}^{j(l')}$.
        Update $\mathbf{X}_k^{j(l')} = \{\mathbf{x}_{1,k}^{j(l')}, \cdots, \mathbf{x}_{n,k}^{j(l')}, \cdots, \mathbf{x}_{N,k}^{j(l')}\}$.
      end
      Set $\mathbf{X}_k^{j(l'+1)} = \mathbf{X}_k^{j(l')}$
    end
    Assign $\mathbf{X}_k^j = \mathbf{X}_k^{j(l)}$ and obtain the pair $\{N_k^j, \widetilde{\mathbf{X}}_k^j\}$.
  end
  `// Selection criterion`
  Select $\{\hat{N}_k, \hat{\mathbf{X}}_k\}$ from the three cases by Equation (20).
  Set the final estimates $\hat{N}_k$ and $\hat{\mathbf{X}}_k$.
  `// Adaptive filtering`
  Adjust I_{k+1} w.r.t. evaluation result from Equation (21).
end

execute the target association according to the RMSE matrix C_k. The targets in the state pairs with the smallest RMSE are then associated together. e_k is then used as a reference criteria to adaptively adjust the particle number and the particle transition range in particle filtering at time $k+1$.

e_k assists adaptive filtering mainly in two aspects: (1) the number of samples I_{k+1} is modified with respect to the RMSE level e_k, a smaller I_{k+1} is assigned when we obtain a lower e_k and vice versa. (2) In the individual dipole propagation, the dipole dynamic range for the sample $\mathbf{x}_k^{ij}$ depends on the value of e_k; a lower e_k results in a smaller dipole dynamic range and vice versa.

4. Numerical results

In this section, we present numerical results using the synthetic data. Since the ground-truth dipole location from real data remains unknown, the performance evaluation relies on the results from the synthetic data.

4.0.1. Simulation setting

We adopted typical examples with both known and unknown number of dipoles. The number of dipoles varied between one and five. The ground-truth dipoles had unit amplitude in our simulation. The orientation of each dipole was set as normal to its corresponding sub-plane. Visualisations were carried out with tools further developed from those published in Helsinki BEM Library [32]. We generated the MEG data using a 204-magnetometer sensor setup. All magnetometers were distributed around the surface of the head. The state space Ω was strictly constrained within the pre-defined 1-layer real head cortex. The width of the brain was 136 mm in our simulation. According to empirical observations, a brain current source often appears and disappears in the same region, and the centre of the current source evolves within a small volume in the cortex. Therefore, it was reasonable for us to assume that all the dipoles were identical and independent of each other, and that each individual dipole might move within a pre-defined triangular sub-plane. We set the measurement SNR (signal to noise ratio) as 10. The measurement noise in the head model had a Gaussian distribution with zero mean and variance σ_ζ^2, where σ_ζ was two times larger than that of the ground-truth noise. We tested each of the algorithms and the model with more than 30 repeated identical experiments. The number of dipoles for each iPF at $k=0$ was set as $I_0 = 10000$.

In the simulation, we test both the GMPF algorithm with one Gibbs iteration (the original GMPF) and the GMPF algorithm with five Gibbs iterations. The GMPF algorithm with five Gibbs iterations outperforms the original GMPF algorithm in unknown dipole number examples, however it only provides slightly better results than that from the original GMPF algorithm in the known dipole number examples. So we only show the results from GMPF algorithm with five Gibbs iterations for the unknown dipole number case in Section 4.2.2. The term 'GMPF' in the rest of this section refers to the original GMPF unless specified otherwise.

In the remainder of this section, we compares the performance of different head models in examples with known and unknown numbers of dipoles. An example with five known numbers of dipoles is used to evaluate the localisation algorithms using the proposed continuous head model. For localisation with an unknown number of dipoles, we test and compare the performance of different particle filtering algorithms in an example with three and dynamic numbers of dipoles. We finally present an evaluation result particularly for the estimation of the dipole number.

4.1. localisation with known number of dipoles

The performance of a known number of dipoles in terms of their localisation is easier and more accurate than the localisation with an unknown number of dipoles. In this section, we present two examples, one with three dipoles (two on the left and one on the right hemisphere) and other with five dipoles (two on the left and three on the right hemisphere). The three dipole example is used to compare three different head models.

4.1.1. Head model comparison

We compared the model performance between the spherical head model, discrete real head model, and the proposed continuous real head model. The spherical head model is a relatively old model that assumes the human head is a perfect spherical shape. This model was used in previous work [14, 13]. The discrete head model is the 1-layer real head model generated using the BEM method, this model was used in [15, 19, 20]. In this paper we adopt the discrete model which contains $G = 8196$ discretised potential source points, all of which have fixed locations. A lead-field matrix with G columns is then generated. We employ the NN interpolation method to convert the discrete model into the continuous model, which was presented in Section 2.

GTM / MM	S/S	S/D	S/C	D/S	D/D	D/C	C/S	C/D	C/C
SIR/2000	17.32	29.37	22.22	35.61	13.01	15.55	21.28	22.78	16.24
MPF/2000	13.33	19.29	19.11	23.65	11.31	13.77	16.34	13.46	**9.22**
GMPF/2000	13.20	19.17	18.16	22.99	11.68	13.94	16.28	12.67	**9.38**
SIR/5000	15.28	28.22	21.35	35.06	11.37	13.10	19.49	17.44	15.36
MPF/5000	11.57	16.32	18.06	21.39	10.73	12.89	16.45	12.08	**7.10**
GMPF/5000	11.77	16.28	18.24	20.12	10.25	12.45	16.99	11.54	**7.53**
SIR/10000	13.05	25.81	18.20	29.77	9.88	11.29	17.53	14.02	12.74
MPF/10000	9.36	15.78	15.13	20.22	9.36	10.97	15.68	9.52	**5.70**
GMPF/10000	9.28	15.90	14.84	20.35	9.33	10.92	15.71	9.48	**5.64**

Table 1: RMSE comparison between three head models using three different particle filter algorithms. All units are in millimeters. GTM: the ground-truth model, MM : the measurement model, S: the spherical head model, D: the discrete head model, C: the continuous head model, SIR: the SIR particle filter, MPF: the multiple particle filter, GMPF: the the proposed Gibbs multiple particle filter, the number in the first column is the sample number employed by the particle filter.

In Table 1, we present the numerical result with three dipoles with known locations. Three models are tested using a simple standard SIR particle filter, a standard MPF and the proposed algorithm. The entries with 'S/S' (column entry) and 'SIR' (row entry) corresponds to the method from [13]; the entries with 'S/S' and 'MPF' is similar to one of the methods adopted in [14]; the entries with 'D/D' and 'MPF' was from our previous work [20].

We vary the particle number from 2000 to 10000. The simulated data were generated and tested using all three models, for example, S/D in the table means that we generate the data using a spherical head model, and the measurement model we use in particle filtering is a discrete head model.

Regarding the RMSE performance between different particle filtering algorithms, we can find that both MPF and GMPF perform better than the standard SIR. MPF and GMPF demonstrated similar performance. This is not surprising since the proposed GMPF algorithm is executed in a similar way to that of MPF in the fixed and known dipole number scenario. We also find that the RMSE performance improves with an increase in particle number.

For those few entries with a similar performance between $I = 2000$ and $I = 5000$ (e.g., column C/S of MPF and GMPF), a similar RMSE may occur due to model mismatching. In terms of head model comparison, it is

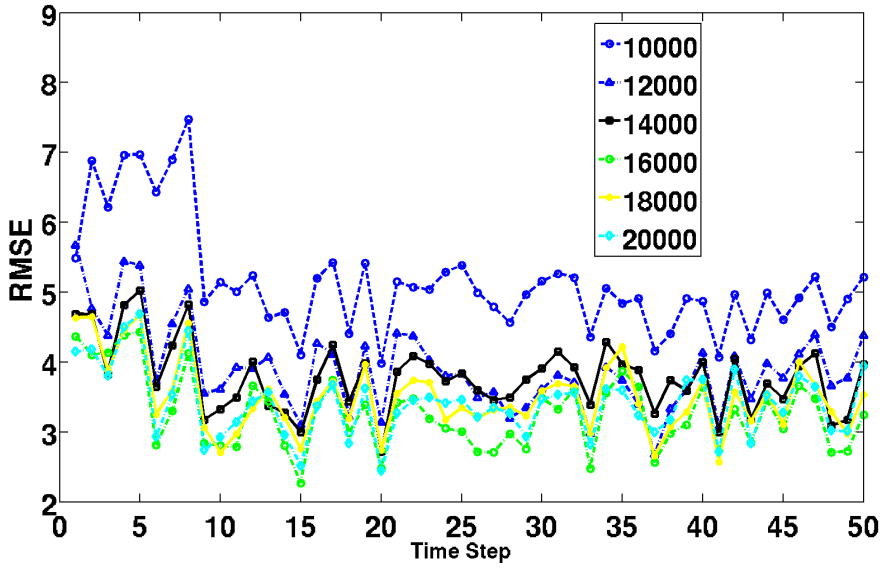

Figure 4: RMSE performance using the C/C model pair and the GMPF algorithm with a varying particle number.

as expected that S/S, D/D, and C/C achieve better performance. Amongst all model pairs for the same algorithm, the continuous head model performs better than the others, as shown in bold. We also find that the spherical model can only perform well when the ground-truth model is the same, while the discrete head model and the continuous head model are more robust to different data.

In Figure 4, we show the detailed RMSE performance using the continuous head model and the proposed GMPF algorithm. For the initial particle numbers larger than 12000, the RMSE stays at the same level. This phenomenon is expected according to the individual dipole dynamic model, as a dipole only moves within the triangular sub-plane.

4.1.2. Known number of five dipoles

Figure 5 shows the tracking results of five known dipoles using the continuous head model and the proposed algorithm.

We can see from Figure 5(a) that the average RMSE over all five dipoles remains at the level of 8 mm. The box-plot ranges between 4 and 12 mm. In Figure 5(b) we show the changes of particle size for some of the identical

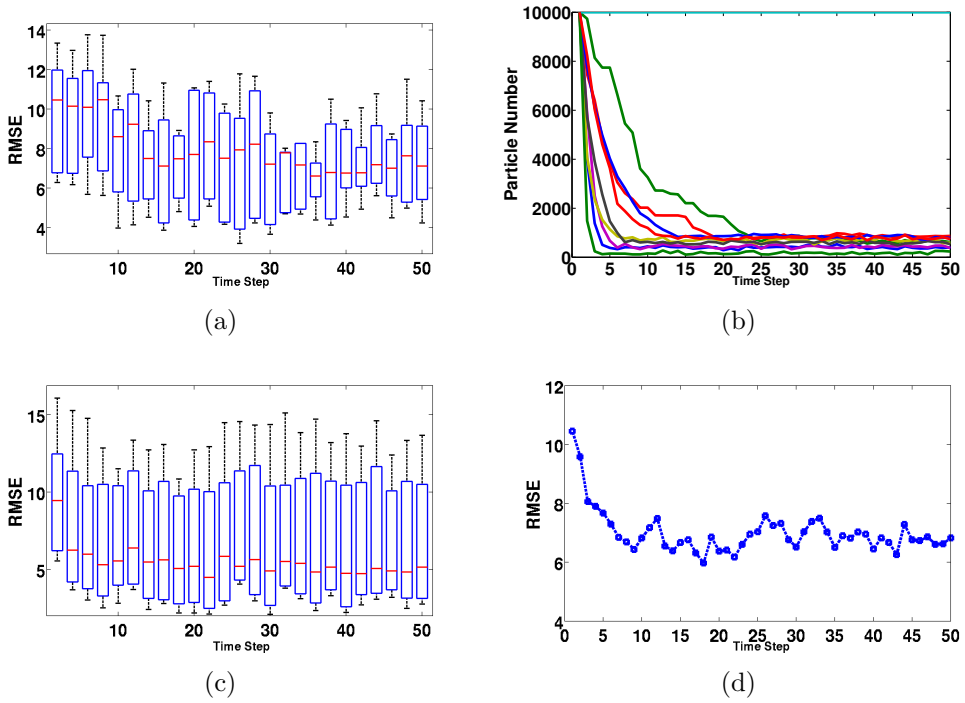

Figure 5: A five-dipole localisation example using the C/C model pair and the GMPF algorithm: (a) shows the average RMSE from the MPF algorithm. (b) shows a decreasing particle size in the algorithm run using the proposed adaptive filtering scheme. The different coloured lines represent the particle size of some identical trails. (c) shows the average RMSE from the GMPF algorithm. (d) shows the RMSE performance of an individual dipole in one of the trails.

trials. The initial particle size is $I_0 = 10000$. We can see that I_k in most experiments dramatically decreases to very low level from around $k = 5$. If we compare the two figures, one can observe that the proposed algorithm is able to achieve a good RMSE performance while adaptively eliminating the number of particles. This greatly saves the computational cost for multiple particle filter types of algorithms. The total particle number is I_k multiplied by the number of dipoles.

From Figure 5 (a) and (c), we can see that the average RMSE over five dipoles is almost at the same level for both MPF and GMPF. This can also

Alg / Dipole	1	2	3	4	5	Processing Time
MPF	8.72	9.85	9.14	10.68	9.57	56 s
GMPF	7.02	9.33	7.58	10.24	9.17	28 s

Table 2: A table comparison for the RMSE and the computer processing time, the particle size for both the MPF and the GMPF algorithms is 10000.

be seen in Table 2. The middle five columns are the dipole index in the five dipole localisation example. This shows the average RMSE for each dipole over 30 identical experiments. We implemented the algorithms in Matlab using a computer with intel Core i7 CPU @ 3.7 GHz. The processing time for an iteration (with 50 time steps) of the MPF algorithm is 56 seconds and that of GMPF is 28 seconds, the total computational time for 30 iterations is around 1820 seconds for the MPF, and 1340 seconds for the proposed method. From the table we observe that the computational time for GMPF is less than that for MPF, this is due to the implementation of the adaptive filtering step in the proposed algorithm. Although we add an extra Gibbs updating step in GMPF, which increases the computational expense), the decreasing of particle number from adaptive filtering step still makes the GMPF faster than the MPF.

4.2. Localisation with unknown number of dipoles

In real world applications, we are not able to obtain prior information of the dipole number, thus an example with an unknown dipole number needs to be tested to access the proposed model and algorithm. The algorithm is tested using two examples: an unknown three dipoles localisation problem and an unknown dynamic dipole numbers localisation. The dipole number in the latter example varies between three and five. The three unknown dipole number examples is used to evaluate the performance of the proposed algorithm.

As we stated in Equation (8), the dipole number dynamics are modelled and estimated through $p(N_k \mid N_{k-1})$. In other words, the new estimated dipole number at time k depends on the previous time step estimates N_{k-1}. For the initial dipole number at time $k = 0$ where there is no historical data available, we assign N_k with the clustered MNE estimation result $\overline{N}_0$.

4.2.1. Evaluation of dipole number estimation

The new dipole number estimates for $k > 0$ highly depend on the estimation at its previous time step. In our selection criterion step, there are three

candidate pairs $\{N_k^j, \mathbf{X}_k^j\}$ where $j = (-), (0), (+)$ as stated previously. In this paper, we multiply the dipole number dynamic probability $p_k(\tau_j)$ from Equation (8) with the probability $p(N_k, \mathbf{X}_k \mid \mathbf{Y}_{1:k})$ computed in the selection criterion step as the candidate posterior probability. We assign $p_k(\tau_0) = 0.5$, $p_k(\tau_-) = 0.25$ and $p_k(\tau_+) = 0.25$. The dipole number dynamic probability can be assigned using other designed schemes. Here we simply give the algorithm a prior knowledge (which comes from the observation using real examples) that the dipole number evolves slowly and it is more likely that the dipole number stays the same between two time steps. N_k is assigned using the candidate pair with the highest normalised candidate posterior probability.

Here we use a three-dipole example to test the proposed approach. The example setting is the same as that in the known dipole example except that no prior dipole number knowledge is given to the algorithm. We apply both the MPF algorithm in paper [20] and the proposed GMPF algorithm to the example using a range of particle numbers. Figure 6 shows the histogram of the dipole number estimates. Each plot is the result from 30 identical experiments (with 50 time step length). Figure 6(a) is the estimation result from the MPF algorithm in paper [20], since its estimation does not depend on the state estimates from the particle filter, we only show one with 10000 particles. We apply our selection criterion assisted GMPF algorithm with a range of particle numbers.

GMPF Particle Number	7000	10000	13000	16000	19000	MPF 10000
Avg. Dipole Number	3.568	3.442	3.260	3.282	3.292	3.562

Table 3: Table shows the mean estimated dipole number using the GMPF algorithm with different particle size. The last column shows the performance from the MPF algorithm with 10000 particles.

Table 3 shows the average estimated dipole number over 30 iterations. Since the ground-truth dipole number is 3, we can observe that with an increase in the particle number from 7000 to 19000, the average estimated dipole number improves slightly. The numbers between GMPF 7000 and GMPF 10000 are similar; however if we check Figure 6(a) and Figure 6(b) for their corresponding histogram distribution, we can see there are more successful estimates (dipole estimate equal to 3) in GMPF 7000 than GMPF 10000. From the histogram we can also observe that, with an increase in the

particle number using GMPF, the result does not improve much for particle sizes larger than 13000.

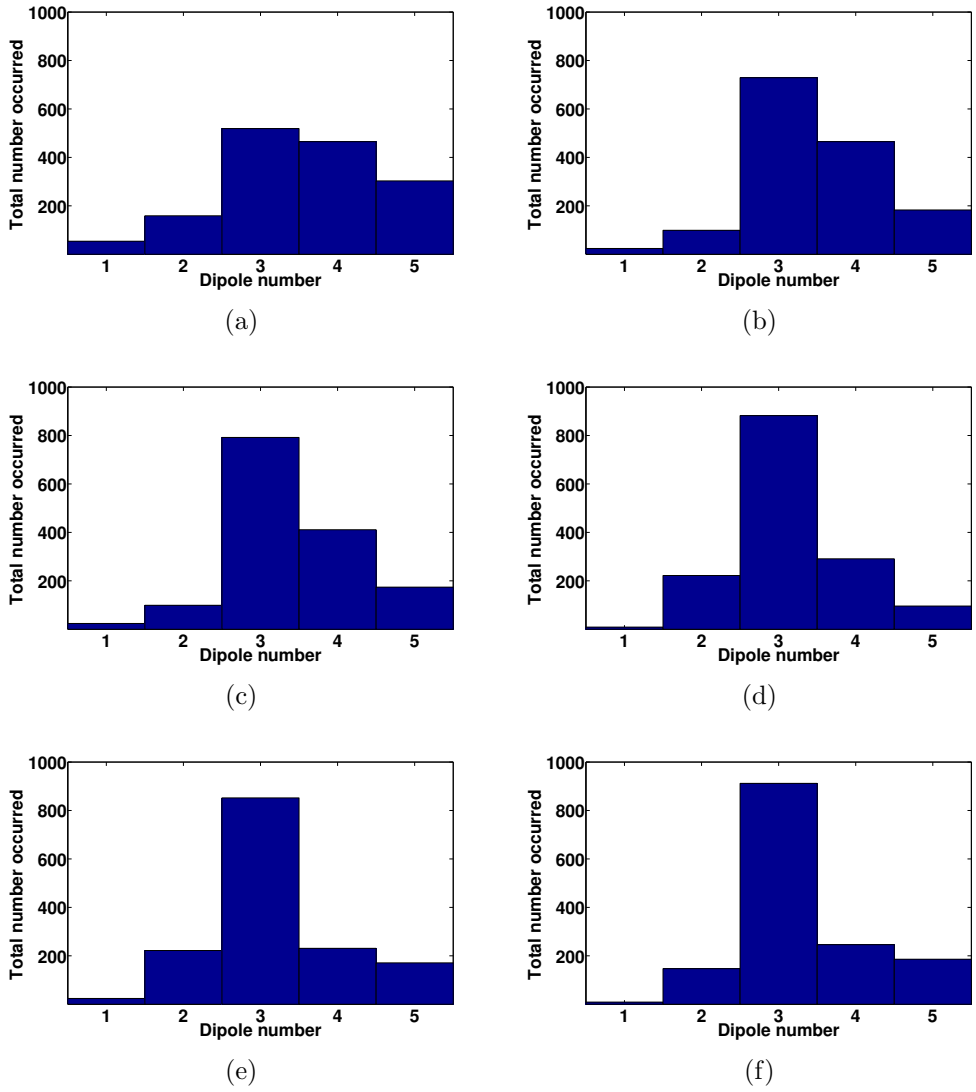

Figure 6: A histogram of the estimated dipole number over 30 identical trails with different particle size: (a) shows the performance using the MPF algorithm with 10000 particles, (b)-(f) shows the performance using the GMPF algorithm with 7000,10000,13000,16000 and 19000 particles, respectively.

4.2.2. Dynamic unknown number of dipole example

We use a dynamic particle number setting in this example. The ground-truth dipole number varies between 3 and 5 during the 50 time steps. The number of dipoles is shown in the 'True' row in Table 4. We use the MPF, the original GMPF and the GMPF with five Gibbs iterations to perform the localisation task; the particle numbers are all equals to 10000.

In order to compare the algorithm performance in multiple target tracking scenario with the unknown number of targets, we adopt a multi-target performance evaluation method based on Optimal Subpattern Assignment Metric (OSPA) method [33]. The method introduces a penalty term to punish the missed or false tracks when there is an unequal estimated dipole number to the ground-truth. We denote the ground-truth dipole number as N_k^v, and the ground-truth dipole state as $\mathbf{X}_k^v$. For case $N_k \geq N_k^G$, the OSPA based distance $D_{p,c}(N_k^v, \mathbf{X}_k^v, N_k, \mathbf{X}_k)$ can be computed by:

$$
\begin{aligned}
&D_{p,c}(N_k^v, \mathbf{X}_k^v, N_k, \mathbf{X}_k) \\
&= [\frac{1}{N_k}(\min_{n_k} \sum_{n_k^v=1}^{N_k^v} (\min(c, \|\mathbf{x}_k - \mathbf{x}_k^v\|_{p'}))^p + (N_k - N_k^v)c^p)]^{\frac{1}{p}},
\end{aligned}
$$

where p and p' are the norm term, c is the penalty term. In this paper, we adopt the value of p = p' = 1 and c = 20, to follow the instructions from [33].

In Figure 7(a) we show the performance of the GMPF with five Gibbs iterations and the plot of the ground-truth dipole number (the blue solid line) versus the average estimated dipole number (the red dashed line) over 50 time steps. The black dashed line is the initial estimate from the probabilistic ROI estimation step. Figure 7(b) shows the corresponding average RMSE for localisation.

We define the average RMSE larger than 30 mm as a lost track, thus the lost track percentages are 16.7%, 20%, 20% for MPF, the original GMPF and the GMPF with five Gibbs iterations, respectively. The following RMSE calculations are based on the results with lost tracks removed. We find that the estimated dipole number follows that of the probabilistic ROI estimation at the first time step.

Table 4 shows the numerical results of the same example for every 5 time steps. Both the estimation of the dipole number and the dipole localisation are acceptable given the standard deviation of the random walk equals to

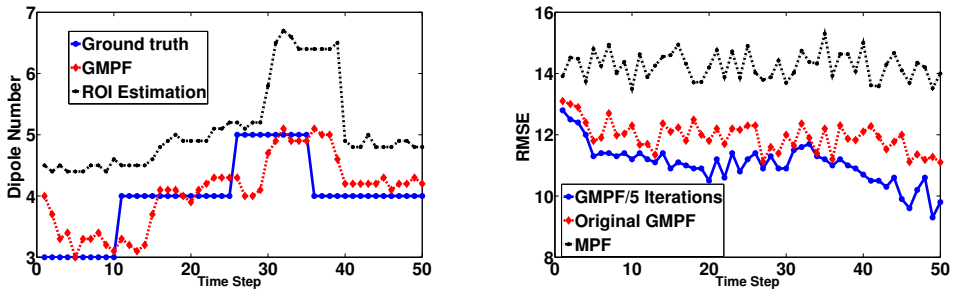

(a) (b)

Figure 7: The estimated dipole number in the dynamic unknown number of dipole example. In (a), the red dashed line shows the mean estimated dipole number using the GMPF algorithm, the blue solid line is the ground-truth dipole number, the black dashed line is the number of the ROIs. (b) shows the localisation RMSE of the corresponding experiments, the blue solid line is the result using a GMPF algorithm with 5 times Gibbs iterations, the red dashed line is the one using a standard GMPF algorithm, and the black dashed line is the one using the MPF algorithm, all with the same particle size.

4 mm. If we compare the results with that achieved by the known dipole number tracking example, the localisation error increases, and the average RMSE is approximately 2 times larger than that in the known dipole number tracking example. This is plausible since we have no prior knowledge of the dipole number. The Gibbs assisted GMPF algorithm has better dipole number estimation performance than the original GMPF with no Gibbs iteration. The RMSE performance for the GMPF with 5 Gibbs iterations is better than that of the original GMPF algorithm, as shown in the table.

5. Conclusion

In this paper, we proposed a continuous real head model and a Bayesian particle filter algorithm to perform multiple dipolar sources localisation in MEG. The proposed algorithm integrated the probabilistic ROI estimation step and the selection criterion step to assist the estimation of the dipole number and the adjustment of the particle size.

The dipole number estimation in the multiple dipole tracking problem

Time k	5	10	15	20	25	30	35	40	45	50
True	3	3	4	4	4	5	5	4	4	4
ROIs	4.4	4.6	4.6	4.9	5.2	5.8	6.4	4.9	4.8	4.8
M/Avg $\hat{N}_k$	3.8	4.1	4.2	4.5	4.7	4.8	5.5	5.0	4.9	4.8
M/RMSE	14.8	13.5	14.6	14.2	14.9	13.7	15.3	15.0	14.1	14.0
O/Avg $\hat{N}_k$	3.5	3.4	4.4	4.3	4.5	4.9	5.6	4.8	4.5	4.4
O/RMSE	11.8	12.3	12.1	11.8	12.3	12.0	12.2	12.1	12.0	11.1
G/Avg $\hat{N}_k$	3	3.1	3.7	3.9	4.3	4.7	4.9	4.2	4.3	4.2
G/RMSE	11.3	11.2	10.9	10.5	11.2	10.9	11.2	10.7	9.9	9.8

Table 4: Table for comparison between the ground-truth dipole number and the estimated dipole number (averaged over 30 iterations). ROIs represents the number of the ROIs estimated using the probabilistic ROI estimation scheme. The M/Avg $\hat{N}_k$ and the M/RMSE represent the mean estimated dipole number and the RMSE performance using the MPF algorithm, respectively. The O/Avg $\hat{N}_k$ and the O/RMSE represent the mean estimated dipole number and the RMSE performance using the original GMPF algorithm with no Gibbs iteration, respectively. The G/Avg $\hat{N}_k$ and the G/RMSE represent the mean estimated dipole number and the RMSE performance using the GMPF algorithm with 5 Gibbs iterations, respectively. All RMSE measures are in unit [mm].

remains a difficult issue. Dimensionality is another issue that needs to be considered. We adopted the GMPF scheme to avoid the dimensionality problem by assigning each target an iPF. We also incorporated the probabilistic ROI estimation step to provide the tracking step with prior information.

However, this method highly depends on the deterministic estimates at the initial step from the MNE. It does not perform well when the dipole number dramatically changes in a short time period. In order to address this problem, we modified the deterministic probabilistic ROI estimation step to a partly probabilistic estimation: we modified the dipole number dynamic model and allowed three potential dipole number guesses. The three candidate pairs are compared and we employed selection criterion scheme to select the pair with the highest probability. We also added Gibbs sampling step to assist tracking in the multiple particle filter. As we have shown in the numerical results, this approach achieves better localisation and dipole number estimation performance than the previously proposed algorithms.

The computational cost of an MPF algorithm approximately equals to

the product of a single particle filter multiplied by the number of identified dipoles. However, it is not necessary to keep a large volume of particles for every time step when the tracking estimate is close enough to the ground-truth. We proposed the adaptive filtering scheme to reduce the computational cost by diminishing particle number when the localisation error is negligible.

We are currently working on a fully probabilistic dipole number estimation method and we will incorporate it into our current localisation algorithm. As the brain current source is a continuous state space in reality, we will explore a better continuous modelling scheme to interpret the brain's current sources. Further research will also focus on decreasing the computational time. This may be addressed by using the parallel computing technique and by implementing a Rao–Blackwelisation method in the linear part of the model.

Acknowledgment

The authors would like to thank Professor Rik Henson and Dr. Olaf Hauk from the MRC Cognition and Brain Sciences Unit in Cambridge, UK, and Dr. Matti Stenroos from Aalto University in Finland, for fruitful discussions and for providing data for the experiments.

6. Reference

References

[1] S. Baillet, J.C. Mosher, and R.M. Leahy, "Electromagnetic brain mapping," *IEEE Signal Processing Magazine*, vol. 18, no. 6, pp. 14–30, 2001.

[2] J. Kaipio and E. Somersalo, *Statistical and computational inverse problems*, vol. 160, Springer, 2005.

[3] M. Hämäläinen, R. Hari, R.J. Ilmoniemi, J. Knuutila, and O.V. Lounasmaa, "Magnetoencephalography: Theory, instrumentation, and applications to noninvasive studies of the working human brain," *Reviews of modern Physics*, vol. 65, no. 2, pp. 413–497, 1993.

[4] Peter Hansen, Morten Kringelbach, and Riitta Salmelin, *MEG: An introduction to methods*, Oxford University Press, 2010.

[5] R.D. Pascual-Marqui, "Standardized low-resolution brain electromagnetic tomography (sLORETA): technical details," *Methods and Findings in Experimental and Clinical Pharmacology*, vol. 24, no. Suppl D, pp. 5–12, 2002.

[6] C.J. Long, P.L. Purdon, S. Temereanca, N.U. Desai, M.S. Hämäläinen, and E.N. Brown, "State-space solutions to the dynamic magnetoencephalography inverse problem using high performance computing," *The annals of applied statistics*, vol. 5, no. 2B, pp. 1207, 2011.

[7] S.C. Jun, J.S. George, J. Paré-Blagoev, S.M. Plis, D.M. Ranken, D.M. Schmidt, and C.C. Wood, "Spatiotemporal Bayesian inference dipole analysis for MEG neuroimaging data," *NeuroImage*, vol. 28, no. 1, pp. 84–98, 2005.

[8] A. Galka, O. Yamashita, T. Ozaki, R. Biscay, and P. Valdés-Sosa, "A solution to the dynamical inverse problem of EEG generation using spatiotemporal Kalman filtering," *NeuroImage*, vol. 23, no. 2, pp. 435–453, 2004.

[9] E. Somersalo, A. Voutilainen, and J.P. Kaipio, "Non-stationary magnetoencephalography by Bayesian filtering of dipole models," *Inverse Problems*, vol. 19, pp. 1047–1063, 2003.

[10] T. Auranen, A. Nummenmaa, M. S. Hämäläinen, I. P. Jääskeläinen, J. Lampinen, A. Vehtari, and M. Sams, "Bayesian analysis of the neuromagnetic inverse problem with ℓ_p-norm priors," *NeuroImage*, vol. 26, no. 3, pp. 870–884, 2005.

[11] A. Nummenmaa, T. Auranen, M. S. Hämäläinen, I. P. Jääskeläinen, J. Lampinen, M. Sams, and A. Vehtari, "Hierarchical Bayesian estimates of distributed MEG sources: theoretical aspects and comparison of variational and MCMC methods," *NeuroImage*, vol. 35, no. 2, pp. 669–685, 2007.

[12] C. Campi, A. Pascarella, A. Sorrentino, and M. Piana, "A Rao–Blackwellized particle filter for magnetoencephalography," *Inverse Problems*, vol. 24, no. 2, pp. 23–37, 2008.

[13] A. Sorrentino, L. Parkkonen, A. Pascarella, C. Campi, and M. Piana, "Dynamical MEG source modeling with multi-target Bayesian filtering," *Human Brain Mapping*, vol. 30, no. 6, pp. 1911–1921, 2009.

[14] L. Miao, J.J. Zhang, C. Chakrabarti, and A. Papandreou-Suppappola, "Efficient Bayesian tracking of multiple sources of neural activity: Algorithms and real-time FPGA implementation," *IEEE Trans on Signal Processing*, vol. 61, no. 3, pp. 633–647, 2013.

[15] A. Sorrentino, A.M. Johansen, J.A.D. Aston, T.E. Nichols, W.S. Kendall, et al., "Dynamic filtering of static dipoles in magnetoencephalography," *The Annals of Applied Statistics*, vol. 7, no. 2, pp. 955–988, 2013.

[16] A.M. Dale, A.K. Liu, B.R. Fischl, R.L. Buckner, J.W. Belliveau, J.D. Lewine, and E. Halgren, "Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity," *Neuron*, vol. 26, no. 1, pp. 55–67, 2000.

[17] J.C. Mosher and R.M. Leahy, "Source localization using recursively applied and projected (RAP) MUSIC," *IEEE Trans. on Signal Processing*, vol. 47, no. 2, pp. 332–340, 1999.

[18] D.M. Schmidt, J.S. George, and C.C. Wood, "Bayesian inference applied to the electromagnetic inverse problem," *Human Brain Mapping*, vol. 7, no. 3, pp. 195–212, 1999.

[19] X. Chen and S.J. Godsill, "Multiple dipolar sources localization for MEG using Bayesian particle filtering," in *Proc. IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2013, pp. 949–953.

[20] X. Chen, S. Sarkka, and S.J. Godsill, "Probabilistic initiation and termination for MEG multiple dipole localization using sequential Monte Carlo methods," in *Proc. 16th Intl. Conf. on Information Fusion (FUSION)*, 2013, pp. 580–587.

[21] G. Nolte, "The magnetic lead field theorem in the quasi-static approximation and its use for magnetoencephalography forward calculation in realistic volume conductors," *Physics in Medicine and Biology*, vol. 48, no. 22, pp. 3637 – 3652, 2003.

[22] M.F. Bugallo, T. Lu, and P.M. Djuric, "Target tracking by multiple particle filtering," in *Proc. IEEE Aerospace Conf.*, 2007, pp. 1–7.

[23] L. Mihaylova, A. Hegyi, A. Gning, and R. Boel, "Parallelized particle and Gaussian sum particle filters for large-scale freeway traffic systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 1, pp. 36–48, 2012.

[24] K.C. Gowda and G. Krishnan, "Agglomerative clustering using the concept of mutual nearest neighbourhood," *Pattern Recognition*, vol. 10, no. 2, pp. 105–112, 1978.

[25] K.J. Friston, A.P. Holmes, K.J. Worsley, J.P. Poline, C.D. Frith, and R.S.J. Frackowiak, "Statistical parametric maps in functional imaging: a general linear approach," *Human brain mapping*, vol. 2, no. 4, pp. 189–210, 1994.

[26] H. S. M. Coxeter, *Introduction to geometry*, Wiley, second edition, 1989.

[27] W. Ng, J. Li, S. Godsill, and S.K. Pang, "Multitarget initiation, tracking and termination using Bayesian Monte Carlo methods," *The Computer Journal*, vol. 50, no. 6, pp. 674–693, 2007.

[28] W.R. Gilks, S. Richardson, and D. Spiegelhalter, *Markov Chain Monte Carlo in Practice*, 1995.

[29] N.J. Gordon, D.J. Salmond, and A.F.M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proc. Part F: Radar and Signal Processing*, vol. 140, no. 2, pp. 107–113, Apr. 1993.

[30] R. Douc and O. Cappe, "Comparison of resampling schemes for particle filtering," in *Proc. Intl. Symp. Image and Signal Processing and Analysis*, Zagreb, Croatia, Sep. 2005, pp. 64–69.

[31] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.

[32] M Stenroos, V Mäntynen, and J Nenonen, "A Matlab library for solving quasi-static volume conduction problems using the boundary element method," *Computer methods and programs in biomedicine*, vol. 88, no. 3, pp. 256–263, 2007.

[33] B. Ristic, B. Vo, D. Clark, and B. Vo, "A metric for performance evaluation of multi-target tracking algorithms," *IEEE Transactions on Signal Processing*, vol. 59, no. 7, pp. 3452–3457, 2011.